arXiv:1704.04315v1 [stat.ME] 14 Apr 2017

Information Criterion for Minimum Cross-Entropy Model Selection

Youngjun Choe
Department of Industrial and Systems Engineering
University of Washington

April 17, 2017

Abstract

The cross-entropy, which is proportional to the Kullback-Leibler divergence, is widely used to gauge the deviation of a distribution of interest from a reference distribution in statistical inference. For example, the Akaike information criterion (AIC) is an asymptotically unbiased estimator of the cross-entropy from a parametric distribution to the true distribution of data. Minimizing the AIC allows us to find a parametric distribution close to the true distribution. In this paper, we generalize the AIC by letting the reference distribution be a target distribution to approximate when its density can be evaluated up to a multiplicative constant only at observed data points. We prove, under some conditions, that the generalized criterion, which we call the cross-entropy information criterion (CIC), is an asymptotically unbiased estimator of the cross-entropy (up to a multiplicative constant) from a parametric distribution to the target distribution. We demonstrate the usefulness of CIC for approximating the optimal importance sampling distribution by a mixture of parametric distributions.

Keywords: cross-entropy information criterion, importance sampling, Kullback-Leibler divergence, mixture model, Monte Carlo

1 Introduction

This paper considers approximating an unknown target density q^* when a nonnegative function r proportional q^* can be evaluated only at observed points. When r is computationally light to evaluate so that discarding some observations is not burdensome, the Metropolis-Hastings algorithm (Metropolis et al., 1953; Hastings, 1970) has been used extensively in Bayesian inference (e.g., q^* is the posterior density and r is the prior density times the likelihood), because of the theoretical guarantee on asymptotically exact sampling from q^*.

But when evaluating r is computationally heavy (e.g., r involves computing the likelihood for a complex Bayesian model (Beaumont, 2010; Sunnåker et al., 2013) or running a computationally intensive algorithm or simulator (Kurtz and Song, 2013; Choe et al., 2015)), we can only afford a small number of the evaluations of r and want to make the best use of each evaluation in approximating q^*. This paper focuses on this scenario where we use every evaluation of r to approximate q^*.

For the approximation, we take a parametric approach by positing a parametric family of densities and finding its member closest to q^* in terms of the closeness measured by the Kullback-Leibler (KL) divergence (Kullback and Leibler, 1951). Thus, the parametric approximation of q^* takes two steps, namely, (1) choosing the parametric family and (2) minimizing the KL divergence for the given family. The latter is a fairly straightforward optimization problem, so the former will be the focus of this paper.

Specifically, the main problem we tackle is how to determine the complexity of the parametric family based on the given data, $(\mathbf{X}_i, r(\mathbf{X}_i)), i = 1, \ldots, n$. If the parametric family is too simple, the parametric approximation of q^* can be poor, whereas if the family is too complex, it can lead to overfitting. We note that this new model selection problem, where r is the only source of information on q^*, is different from the traditional model selection problem where the key information is contained in data, $\mathbf{X}_i, i = 1, \ldots, n$, directly sampled from an unknown density to model. The traditional problem is extensively studied (Burnham and Anderson, 2003) and its most well-known model selection criteria include the Akaike information criterion (AIC) (Akaike, 1973, 1974) and Bayesian information criterion (BIC) (Schwarz, 1978).

To the best of our knowledge, the new model selection problem is not yet investigated despite its potential importance in Bayesian inference (Beaumont, 2010; Sunnåker et al., 2013) and Monte Carlo studies (Kurtz and Song, 2013; Choe et al., 2015). To address the problem, we propose a novel information criterion and prove that the criterion is an asymptotically unbiased estimator of the KL divergence (up to a multiplicative constant and an additive constant), justifying the minimization of the criterion for selecting the best model.

The remainder of this paper is organized as follows. Section 2 provides the background. Section 3 proposes the novel information criterion. Section 4 explains how the proposed criterion can be used in practice to adaptively posit an increasingly complex parametric family to approximate an unknown target density q^* as more evaluations of r proportional to q^* are amassed. Section 5 numerically demonstrates the use of the proposed criterion for approximating the optimal density for importance sampling (Kahn and Marshall, 1953). Section 6 concludes and suggests future research directions.

2 Background

The KL divergence (Kullback and Leibler, 1951) is commonly used to gauge the difference between two distributions in statistical inference. For two probability measures, Q^* and Q, on a common measurable space such that Q^* is absolutely continuous with respect to Q (written $Q^* \ll Q$), the KL divergence from Q to Q^* is

$$\mathbb{D}(Q^*||Q) := \mathbb{E}_{Q^*}\left[\log \frac{\mathrm{d}Q^*}{\mathrm{d}Q}\right], \tag{1}$$

where $\mathbb{E}_{Q^*}$ denotes the expectation with respect to Q^* and $\mathrm{d}Q^*/\mathrm{d}Q$ is the Radon-Nikodym derivative of Q^* with respect to Q. If Q^* and Q are absolutely continuous with respect to a dominating measure μ (e.g., counting or Lebesgue measure), their respective densities, q^* and q, exist by the Radon-Nikodym theorem and the KL

divergence in (1) can be expressed as

$$\begin{aligned}\mathbb{D}(Q^*||Q) &= \mathbb{E}_{Q^*}\left[\log \frac{q^*}{q}\right] \\ &= \int q^* \log\left(\frac{q^*}{q}\right) \mathrm{d}\mu.\end{aligned}$$

The KL divergence is non-negative and takes zero if and only if $q^* = q$ almost everywhere (a.e.). If Q is an approximation of Q^*, the choice of Q minimizing $\mathbb{D}(Q^*||Q)$ leads to a good approximation.

The maximum likelihood estimator (MLE) is a prominent example of using the KL divergence. When a sample of p–dimensional random vectors $\mathbf{X}_1, \ldots, \mathbf{X}_n$ is drawn from an unknown distribution Q^*, we can approximate Q^* by a distribution in a parametric family $\{Q_{\boldsymbol{\theta}} : \boldsymbol{\theta} \in \boldsymbol{\Theta} \subset \mathbb{R}^d\}$ by minimizing the KL divergence from $Q_{\boldsymbol{\theta}}$ to Q^* over $\boldsymbol{\theta} \in \boldsymbol{\Theta}$. Suppose $Q^* \ll Q_{\boldsymbol{\theta}}$ for all $\boldsymbol{\theta} \in \boldsymbol{\Theta}$ so that the KL divergence is well defined over $\boldsymbol{\Theta}$. Also, suppose $Q_{\boldsymbol{\theta}} \ll \mu$ for all $\boldsymbol{\theta} \in \boldsymbol{\Theta}$ and $Q^* \ll \mu$ so that densities $q_{\boldsymbol{\theta}} = \mathrm{d}Q_{\boldsymbol{\theta}}/\mathrm{d}\mu$ and $q^* = \mathrm{d}Q^*/\mathrm{d}\mu$ exist. Then, the KL divergence is

$$\begin{aligned}\mathbb{D}(Q^*||Q_{\boldsymbol{\theta}}) &= \int q^* \log\left(\frac{q^*}{q_{\boldsymbol{\theta}}}\right) \mathrm{d}\mu \\ &= \int q^* \log q^* \, \mathrm{d}\mu - \int q^* \log q_{\boldsymbol{\theta}} \, \mathrm{d}\mu. \qquad (2)\end{aligned}$$

Note that only the second term in (2), called cross-entropy, depends on $\boldsymbol{\theta}$. Therefore, minimizing the KL divergence over $\boldsymbol{\theta} \in \boldsymbol{\Theta}$ is equivalent to minimizing the cross-entropy over $\boldsymbol{\theta} \in \boldsymbol{\Theta}$. Because q^* is unknown, the cross-entropy should be estimated based on $\mathbf{X}_1, \ldots, \mathbf{X}_n \sim Q^*$. An unbiased, consistent estimator of the cross-entropy is

$$-\frac{1}{n}\sum_{i=1}^{n} \log q_{\boldsymbol{\theta}}(\mathbf{X}_i), \qquad (3)$$

which is the average of negative log-likelihoods for the observed data. Therefore, the MLE of $\boldsymbol{\theta}$, denoted by $\widehat{\boldsymbol{\theta}}_n$, is the minimizer of the cross-entropy estimator in (3).

Another example of using the KL divergence or cross-entropy is the AIC (Akaike, 1973, 1974). As the free parameter dimension d of the parameter space $\boldsymbol{\Theta}$ (or equivalently, the model degrees of freedom) increases, $Q_{\widehat{\boldsymbol{\theta}}_n}$ may become a better approxi-

mation of Q^*. To compare the different approximating distributions (or models), we could use a plug-in estimator

$$-\frac{1}{n}\sum_{i=1}^{n}\log q_{\widehat{\boldsymbol{\theta}}_n}(\mathbf{X}_i) \tag{4}$$

of the cross-entropy, but this is problematic because of the downward bias created from using the data twice. While complex models (or distributions with larger d's) tend to have a smaller cross-entropy estimate, they are subject to the overfitting problem. The AIC remedies this issue by correcting the asymptotic bias of the estimator in (4). The AIC is defined (up to a multiplicative constant) as

$$-\frac{1}{n}\sum_{i=1}^{n}\log q_{\widehat{\boldsymbol{\theta}}_n}(\mathbf{X}_i)+\frac{d}{n}, \tag{5}$$

where the bias correction term d/n penalizes the model complexity, balancing it with the goodness-of-fit represented by the first term. Minimizing the AIC can be interpreted as minimizing an asymptotically unbiased estimator of the cross-entropy. Therefore, both the MLE and AIC aim at minimizing the cross-entropy from an approximate distribution $Q_{\boldsymbol{\theta}}$ to the unknown distribution Q^* generating the data.

Analogous to the MLE and AIC, our approximation task in this paper considers minimizing the KL divergence (or cross-entropy) from a parametric distribution $Q_{\boldsymbol{\theta}}$ to the target distribution Q^* over $\boldsymbol{\theta}\in\boldsymbol{\Theta}$ (to well-define the KL divergence, hereafter assume $Q^*\ll Q_{\boldsymbol{\theta}}$ for all $\boldsymbol{\theta}\in\boldsymbol{\Theta}$) when the target density q^* is proportional to a non-negative function r (i.e., $q^*=r/\rho$ for a positive constant $\rho=\int r\,\mathrm{d}\mu$). Minimizing the KL divergence in (2) (equivalently, cross-entropy) over $\boldsymbol{\theta}$ is equivalent to minimizing

$$\mathcal{C}(\boldsymbol{\theta}):=-\int r\log q_{\boldsymbol{\theta}}\,\mathrm{d}\mu, \tag{6}$$

which is unknown in practice because r can be evaluated only at observed data points. We can approximate $\mathcal{C}(\boldsymbol{\theta})$ in (6) by an unbiased, consistent estimator

$$\bar{\mathcal{C}}_{\boldsymbol{\eta}}(\boldsymbol{\theta}):=-\frac{1}{n}\sum_{i=1}^{n}\frac{r(\mathbf{X}_i)}{q_{\boldsymbol{\eta}}(\mathbf{X}_i)}\log q_{\boldsymbol{\theta}}(\mathbf{X}_i) \tag{7}$$

for $\mathbf{X}_1, \ldots, \mathbf{X}_n \sim Q_{\boldsymbol{\eta}}$ with any parameter $\boldsymbol{\eta} \in \boldsymbol{\Theta}$, because $\mathbb{E}_{Q_{\boldsymbol{\eta}}} \bar{\mathcal{C}}_{\boldsymbol{\eta}}(\boldsymbol{\theta}) = \mathcal{C}(\boldsymbol{\theta})$. Therefore, by minimizing $\bar{\mathcal{C}}_{\boldsymbol{\eta}}(\boldsymbol{\theta})$ in (7) over $\boldsymbol{\theta} \in \boldsymbol{\Theta}$, we can approximately minimize the KL divergence (or cross-entropy) from $Q_{\boldsymbol{\theta}}$ to Q^*. Thus, we call $\widehat{\boldsymbol{\theta}}_n := \operatorname{argmin}_{\boldsymbol{\theta} \in \boldsymbol{\Theta}} \bar{\mathcal{C}}_{\boldsymbol{\eta}}(\boldsymbol{\theta})$ the minimum cross-entropy estimator (MCE). Note that if the random sample is drawn from the target distribution (i.e., $\mathbf{X}_1, \ldots, \mathbf{X}_n \sim Q_{\boldsymbol{\eta}} = Q^*$), then the MCE reduces to the MLE because minimizing (7) is equivalent to minimizing (3) due to $r/q^* = \rho$.

3 Cross-Entropy Information Criterion

This section develops the cross-entropy information criterion (CIC) that is to the AIC as the MCE is to the MLE. For a fixed parametric family $\{Q_{\boldsymbol{\theta}} : \boldsymbol{\theta} \in \boldsymbol{\Theta}\}$, the MCE $\widehat{\boldsymbol{\theta}}_n$ yields $Q_{\widehat{\boldsymbol{\theta}}_n}$ closest to Q^* with respect to the data. The CIC will guide us to select the dimension d of $\boldsymbol{\Theta}$, allowing us to compare different approximating distributions (or models).

As the AIC balances the goodness-of-fit and the model complexity, the CIC should balance the cross-entropy estimate and the model complexity. We derive the model complexity penalty term for the CIC by following a similar path for deriving the AIC, i.e., we derive the downward bias introduced from estimating $\mathcal{C}(\widehat{\boldsymbol{\theta}}_n)$ by $\bar{\mathcal{C}}_{\boldsymbol{\eta}}(\widehat{\boldsymbol{\theta}}_n)$, where $\mathcal{C}(\widehat{\boldsymbol{\theta}}_n)$ is the true cross-entropy (up to a multiplicative constant) from $Q_{\widehat{\boldsymbol{\theta}}_n}$ to Q^*. To derive the bias in a closed-form, we need regularity conditions similar to those needed for deriving the AIC (see Assumptions 2–16 in Appendix A).

3.1 Properties of the minimum cross-entropy estimator

Because the CIC is based on the MCE, we first characterize the asymptotic behavior of the MCE. The MCE, which is a minimizer of (7), is an M-estimator (Huber, 1964), so that the M-estimation theory (Van der Vaart, 1998) applies to the MCE. Specifically, assume the parameter minimizing the true cross-entropy, $\boldsymbol{\theta}^* := \operatorname{argmin}_{\boldsymbol{\theta} \in \boldsymbol{\Theta}} \mathcal{C}(\boldsymbol{\theta})$, is unique hereafter. Then, MCE $\widehat{\boldsymbol{\theta}}_n$ is a strongly consistent estimator of $\boldsymbol{\theta}^*$ (Lemma 1) and is asymptotically normal (Lemma 2) under standard regularity conditions (see Assumptions 2–8 in Appendix A).

Lemma 1 (Strong consistency of the MCE)**.** *Suppose that Assumptions 2–3 hold and* $\boldsymbol{\Theta}$ *is compact. Then, for any* $\boldsymbol{\eta} \in \boldsymbol{\Theta}$*,*

$$\widehat{\boldsymbol{\theta}}_n := \operatorname*{argmin}_{\boldsymbol{\theta} \in \boldsymbol{\Theta}} \bar{\mathcal{C}}_{\boldsymbol{\eta}}(\boldsymbol{\theta})$$

converges almost surely to $\boldsymbol{\theta}^*$ *as* $n \to \infty$*.*

Proof. *This is a direct result of Theorem A1 in Rubinstein and Shapiro (1993).*

Lemma 2 (Asymptotic normality of the MCE)**.** *Suppose that Assumptions 4–8 hold and that for any* $\boldsymbol{\eta} \in \boldsymbol{\Theta}$*,* $\widehat{\boldsymbol{\theta}}_n$ *converges in probability to* $\boldsymbol{\theta}^*$ *as* $n \to \infty$*. Then, for any* $\boldsymbol{\eta} \in \boldsymbol{\Theta}$*,* $\sqrt{n}\left(\widehat{\boldsymbol{\theta}}_n - \boldsymbol{\theta}^*\right)$ *converges in distribution to* $N\left(0, \boldsymbol{\Gamma}^{-1}\boldsymbol{\Lambda}_{\boldsymbol{\eta}}\boldsymbol{\Gamma}^{-1}\right)$ *as* $n \to \infty$*, where*

$$\boldsymbol{\Gamma} := -\mathbb{E}_{\mu}\left[r \nabla^2_{\boldsymbol{\theta}} \log q_{\boldsymbol{\theta}^*}\right]$$

and

$$\boldsymbol{\Lambda}_{\boldsymbol{\eta}} := \mathbb{E}_{Q_{\boldsymbol{\eta}}}\left[\nabla_{\boldsymbol{\theta}} h(\mathbf{X}, \boldsymbol{\eta}, \boldsymbol{\theta}^*) \nabla_{\boldsymbol{\theta}} h(\mathbf{X}, \boldsymbol{\eta}, \boldsymbol{\theta}^*)^T\right]$$

for

$$h(\mathbf{X}, \boldsymbol{\eta}, \boldsymbol{\theta}) := \frac{r(\mathbf{X})}{q_{\boldsymbol{\eta}}(\mathbf{X})} \log q_{\boldsymbol{\theta}}(\mathbf{X}). \tag{8}$$

Proof. *This is a direct result of Theorem A2 in Rubinstein and Shapiro (1993).*

3.2 Iterative procedure for minimizing the cross-entropy

In contrast to the AIC that uses the data drawn from the target distribution to approximate, the CIC uses the data drawn from a distribution (e.g., $Q_{\boldsymbol{\eta}}$) different from Q^*. To compensate for this lack of information, the data distribution needs to asymptotically approach Q^*. Therefore, we consider an iterative procedure which refines the data distribution as we gather data. Specifically, the structure of the estimator in (7) inspires the iterative procedure in Figure 1 with the maximum number

of iterations, $\tau \in [2, \infty)$. We note that a special case of this procedure is known as the cross-entropy method (Rubinstein, 1999) when r in $\bar{\mathcal{C}}_{\boldsymbol{\eta}}(\boldsymbol{\theta})$ of Step 2 is proportional to the optimal importance sampling density.

ITERATIVE PROCEDURE FOR APPROXIMATING Q^*

Inputs: iteration counter $t = 1$, the maximum iteration τ, and the initial parameter $\widehat{\boldsymbol{\theta}}_n^{(0)} = \boldsymbol{\eta} \in \boldsymbol{\Theta}$.

1. Sample $\mathbf{X}_1^{(t-1)}, \ldots, \mathbf{X}_n^{(t-1)}$ from $Q_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}$.
2. Find the MCE $\widehat{\boldsymbol{\theta}}_n^{(t)} := \underset{\boldsymbol{\theta} \in \boldsymbol{\Theta}}{\operatorname{argmin}} \bar{\mathcal{C}}_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}(\boldsymbol{\theta})$, where

$$\begin{aligned} \bar{\mathcal{C}}_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}(\boldsymbol{\theta}) &:= -\frac{1}{n} \sum_{i=1}^{n} \frac{r\left(\mathbf{X}_i^{(t-1)}\right)}{q_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}\left(\mathbf{X}_i^{(t-1)}\right)} \log q_{\boldsymbol{\theta}}\left(\mathbf{X}_i^{(t-1)}\right), \\ &= -\frac{1}{n} \sum_{i=1}^{n} h\left(\mathbf{X}_i^{(t-1)}, \widehat{\boldsymbol{\theta}}_n^{(t-1)}, \boldsymbol{\theta}\right) \end{aligned} \tag{9}$$

3. If $t = \tau$, output the approximate distribution $Q_{\widehat{\boldsymbol{\theta}}_n^{(\tau)}}$. Otherwise, increment t by 1 and go to Step 1.

Figure 1: Iterative procedure for approximating Q^* by minimizing the estimator of cross-entropy from a parametric distribution $Q_{\boldsymbol{\theta}}$ to Q^*.

We also note that $\bar{\mathcal{C}}_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}(\boldsymbol{\theta})$ in (9) is an unbiased, consistent estimator of the cross-entropy (up to a multiplicative constant) from $Q_{\boldsymbol{\theta}}$ to Q^*, making $\widehat{\boldsymbol{\theta}}_n^{(t)} := \underset{\boldsymbol{\theta} \in \boldsymbol{\Theta}}{\operatorname{argmin}} \bar{\mathcal{C}}_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}(\boldsymbol{\theta})$ the MCE. We use the same sample size n for each iteration for notational simplicity without loss of generality.

In this algorithm, the MCE $\widehat{\boldsymbol{\theta}}_n^{(t)}$ in each iteration is strongly consistent as stated in Corollary 3. We use this consistency and the iterative procedure to establish our main theoretical results justifying the CIC in the next section.

Corollary 3. *Suppose the conditions in Lemma 1 hold. Then,*

$$\widehat{\boldsymbol{\theta}}_n^{(t)} := \underset{\boldsymbol{\theta} \in \boldsymbol{\Theta}}{\operatorname{argmin}} \bar{\mathcal{C}}_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}(\boldsymbol{\theta}) \tag{10}$$

converges almost surely to $\boldsymbol{\theta}^*$ *as* $n \to \infty$ *for* $t = 1, \ldots, \tau$.

Proof. $\widehat{\boldsymbol{\theta}}_n^{(t-1)} \in \boldsymbol{\Theta}$ *holds for* $n \geq 1$ *and* $t = 1, \ldots, \tau$. *The convergence follows from Lemma 1.*

3.3 Asymptotic bias of a cross-entropy estimator

To simplify the model complexity penalty term in the AIC, Akaike (1974) makes the strong assumption that the true distribution of data belongs to the parametric distribution family being considered. We similarly make Assumption 1 to simplify the asymptotic bias of $\bar{\mathcal{C}}_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}\left(\widehat{\boldsymbol{\theta}}_n^{(t)}\right)$ in estimating $\mathcal{C}\left(\widehat{\boldsymbol{\theta}}_n^{(t)}\right)$, because the asymptotic bias corresponds to the penalty term in the CIC.

Assumption 1. *There exists* $\boldsymbol{\theta}^* \in \boldsymbol{\Theta}$ *such that* $Q^* = Q_{\boldsymbol{\theta}^*}$.

This assumption implies that the target distribution Q^* belongs to the parametric family $\{Q_{\boldsymbol{\theta}} : \boldsymbol{\theta} \in \boldsymbol{\Theta}\}$ in which the approximating distribution is searched. Under Assumption 1, we can establish the asymptotic normality of the MCE $\widehat{\boldsymbol{\theta}}_n^{(t)}$ in Corollary 4. This leads to our main result in Theorem 5 quantifying the asymptotic bias (see Appendix A for the proofs of Corollary 4 and Theorem 5).

Corollary 4. *Suppose that Assumptions 1 and 4–10 hold and that* $\widehat{\boldsymbol{\theta}}_n^{(t)} := \operatorname{argmin}_{\boldsymbol{\theta} \in \boldsymbol{\Theta}} \bar{\mathcal{C}}_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}(\boldsymbol{\theta})$ *converges in probability to* $\boldsymbol{\theta}^*$ *as* $n \to \infty$ *for* $t = 1, \ldots, \tau$. *Then,* $\boldsymbol{\Lambda}_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}$ *converges in probability to* $\rho\boldsymbol{\Gamma}$ *and* $\sqrt{n}\left(\widehat{\boldsymbol{\theta}}_n^{(t)} - \boldsymbol{\theta}^*\right)$ *converges in distribution to* $N\left(0, \rho\boldsymbol{\Gamma}^{-1}\right)$ *as* $n \to \infty$ *for* $t = 2, \ldots, \tau$, *where* $\rho = \mathbb{E}_\mu[r]$.

Theorem 5 (Asymptotic bias of $\bar{\mathcal{C}}_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}\left(\widehat{\boldsymbol{\theta}}_n^{(t)}\right)$ in estimating $\mathcal{C}\left(\widehat{\boldsymbol{\theta}}_n^{(t)}\right)$)**.** *Suppose that Assumptions 1, 5–6, and 11–16 hold, that* $\widehat{\boldsymbol{\theta}}_n^{(t)}$ *converges in probability to* $\boldsymbol{\theta}^*$ *as* $n \to \infty$ *for* $t = 1, \ldots, \tau$, *and that* $\sqrt{n}\left(\widehat{\boldsymbol{\theta}}_n^{(t)} - \boldsymbol{\theta}^*\right)$ *converges in distribution to* $N\left(0, \rho\boldsymbol{\Gamma}^{-1}\right)$ *as* $n \to \infty$ *for* $t = 2, \ldots, \tau$. *Then,*

$$\mathbb{E}\left[\bar{\mathcal{C}}_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}\left(\widehat{\boldsymbol{\theta}}_n^{(t)}\right) - \mathcal{C}\left(\widehat{\boldsymbol{\theta}}_n^{(t)}\right)\right] = -\rho\frac{d}{n} + o\left(\frac{1}{n}\right)$$

for $t = 2, \ldots, \tau$.

The asymptotic bias, $-\rho d/n$, is proportional to the free parameter dimension d of the parameter space $\boldsymbol{\Theta}$, similar to the penalty term of the AIC in (5). In practice,

$\rho = \mathbb{E}_\mu[r]$ is unknown, but we can use a consistent estimator of ρ to estimate the asymptotic bias. For example, at the t^{th} iteration, $t = 1, \ldots, \tau$, we can use

$$\widehat{\rho}^{(t-1)} = \frac{1}{n} \sum_{i=1}^{n} \frac{r\left(\mathbf{X}_i^{(t-1)}\right)}{q_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}\left(\mathbf{X}_i^{(t-1)}\right)} \tag{11}$$

as an unbiased, consistent estimator of ρ. Furthermore, if $q_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}} = q^*$, the estimator in (11) is the optimal importance sampling estimator having zero variance (Kahn and Marshall, 1953). Because the iterative procedure refines $q_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}$ to be closer to q^*, it is generally expected that $\widehat{\rho}^{(t-1)}$ will have a smaller variance as t gets larger.

The fact that we can estimate ρ with a small variance as a byproduct of approximating q^* is particularly desirable, because ρ, the normalizing constant of $q^* = r/\rho$, is often a quantity of interest. For example, if r is the prior density times the likelihood of a model, then ρ is the evidence for the model, which can be compared between models for Bayesian model selection (Knuth et al., 2015).

Because ρ can be consistently estimated, we can compute the CIC, which is a bias-corrected estimator of $\mathcal{C}\left(\widehat{\boldsymbol{\theta}}_n^{(t)}\right)$, the cross-entropy up to a multiplicative constant:

Cross-entropy information criterion (CIC)

$$\mathrm{CIC}^{(t)}(d) = \bar{\mathcal{C}}_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}\left(\widehat{\boldsymbol{\theta}}_n^{(t)}\right) + \widehat{\rho}\frac{d}{n} \tag{12}$$

for $t = 1, \ldots, \tau$, where $\widehat{\boldsymbol{\theta}}_n^{(t)} := \underset{\boldsymbol{\theta} \in \boldsymbol{\Theta}}{\operatorname{argmin}} \bar{\mathcal{C}}_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}(\boldsymbol{\theta})$ with $\bar{\mathcal{C}}_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}(\cdot)$ in (9). $\widehat{\rho}$ is a consistent estimator of ρ, such as the estimator in (11).

We note that the parameter space $\boldsymbol{\Theta}$, in which the MCE $\widehat{\boldsymbol{\theta}}_n^{(t)}$ is found, is a function of the free parameter dimension d, although it is not explicitly expressed for notational simplicity. For example, if the parametric family is a mixture of k Gaussian component distributions parametrized by their means and covariances, then the parameter space $\boldsymbol{\Theta}$ is fully specified once we select the number of components, k, which directly determines the free parameter dimension d (see Section 4.2 for more details on this example). Therefore, for a fixed d, we can find the MCE $\widehat{\boldsymbol{\theta}}_n^{(t)}$ and then compute $\mathrm{CIC}^{(t)}(d)$ in (12).

We note that the CIC reduces to the AIC up to an additive $o_p(1/n)$ if the samples are all drawn from the target distribution, that is, $\mathbf{X}_1^{(t-1)}, \ldots, \mathbf{X}_n^{(t-1)} \sim Q_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}} = Q^*$ for $t = 1, \ldots, \tau$ in Figure 1. If so, the first term of the CIC in (12) becomes

$$\bar{\mathcal{C}}_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}\left(\widehat{\boldsymbol{\theta}}_n^{(t)}\right) := -\frac{1}{n}\sum_{i=1}^{n} \frac{r\left(\mathbf{X}_i^{(t-1)}\right)}{q_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}\left(\mathbf{X}_i^{(t-1)}\right)} \log q_{\widehat{\boldsymbol{\theta}}_n^{(t)}}\left(\mathbf{X}_i^{(t-1)}\right) \tag{13}$$

$$= -\frac{1}{n}\sum_{i=1}^{n} \frac{r\left(\mathbf{X}_i^{(t-1)}\right)}{q^*\left(\mathbf{X}_i^{(t-1)}\right)} \log q_{\widehat{\boldsymbol{\theta}}_n^{(t)}}\left(\mathbf{X}_i^{(t-1)}\right) \tag{14}$$

$$= -\frac{\rho}{n}\sum_{i=1}^{n} \log q_{\widehat{\boldsymbol{\theta}}_n^{(t)}}\left(\mathbf{X}_i^{(t-1)}\right), \tag{15}$$

because $q_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}} = q^*$ in (13) and $r/q^* = \rho$ in (14). By plugging the expression in (15) into the CIC in (12), it shows that the CIC is equal to ρ times the AIC in (5) up to an additive $o_p(1/n)$.

The asymptotic bias expression in Theorem 5 holds only for $t \geq 2$, because when $t = 1$, the initial sample is drawn from $Q_{\widehat{\boldsymbol{\theta}}_n^{(0)}} = Q_{\boldsymbol{\eta}}$, which is not a distribution converging to $Q_{\boldsymbol{\theta}^*}$. If we want to select a reasonable parameter dimension d to use at the first iteration, it is still necessary to penalize the model complexity. Therefore, we define the CIC even for $t = 1$.

If we use the same sample size n for each iteration, the model dimension d for later iterations may vary only a little from the earlier iterations. Alternatively, we can allocate larger n to later iterations, or aggregate the samples gathered through iterations to obtain a cumulative version of the CIC as proposed in Section 4. The potential benefit of the latter approach is the tendency of the aggregated estimator of $\mathcal{C}(\boldsymbol{\theta})$ to have a smaller variance than the non-aggregated estimator $\bar{\mathcal{C}}_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}(\boldsymbol{\theta})$ in (9). Section 4 gives the details.

4 Application of the Cross-Entropy Information Criterion

4.1 The CIC based on cumulative data

To reduce the variance of the MCE $\widehat{\boldsymbol{\theta}}_n^{(t)} := \underset{\boldsymbol{\theta}\in\boldsymbol{\Theta}}{\operatorname{argmin}}\, \bar{\mathcal{C}}_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}(\boldsymbol{\theta})$, instead of using only the last iteration's data $\mathbf{X}_1^{(t-1)}, \ldots, \mathbf{X}_n^{(t-1)} \sim Q_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}$, we can use all the observed data up to the last iteration to estimate $\mathcal{C}(\boldsymbol{\theta})$. For more flexibility, we can allocate a different sample size for each iteration, i.e., n_t for the t^{th} iteration, $t = 0, 1, \ldots, \tau$ (e.g., a large n_0 for the initial sample to broadly cover the support of $Q_{\boldsymbol{\eta}}$ and equal sample sizes $n_1 = \ldots = n_\tau$ for the later iterations). Then, we can find the MCE

$$\widehat{\boldsymbol{\theta}}^{(t)} := \underset{\boldsymbol{\theta}\in\boldsymbol{\Theta}}{\operatorname{argmin}}\, \bar{\mathcal{C}}^{(t-1)}(\boldsymbol{\theta}), \tag{16}$$

where

$$\begin{aligned}
\bar{\mathcal{C}}^{(t-1)}(\boldsymbol{\theta}) &:= \frac{1}{\sum_{s=0}^{t-1} n_s} \sum_{s=0}^{t-1} n_s \bar{\mathcal{C}}_{\widehat{\boldsymbol{\theta}}^{(s)}}(\boldsymbol{\theta}) \\
&= -\frac{1}{\sum_{s=0}^{t-1} n_s} \sum_{s=0}^{t-1} \sum_{i=1}^{n_s} \frac{r\left(\mathbf{X}_i^{(s)}\right)}{q_{\widehat{\boldsymbol{\theta}}^{(s)}}\left(\mathbf{X}_i^{(s)}\right)} \log q_{\boldsymbol{\theta}}\left(\mathbf{X}_i^{(s)}\right) \\
&= -\frac{1}{\sum_{s=0}^{t-1} n_s} \sum_{s=0}^{t-1} \sum_{i=1}^{n_s} h\left(\mathbf{X}_i^{(s)}, \widehat{\boldsymbol{\theta}}^{(s)}, \boldsymbol{\theta}\right)
\end{aligned} \tag{17}$$

for $t = 1, \ldots, \tau$ with $\widehat{\boldsymbol{\theta}}^{(0)} := \boldsymbol{\eta}$. We note that $\bar{\mathcal{C}}^{(t-1)}(\boldsymbol{\theta})$ in (17) is an unbiased estimator of $\mathcal{C}(\boldsymbol{\theta})$.

By using all data gathered until the t^{th} iteration, we can determine the model parameter dimension d at the t^{th} iteration with the following CIC:

Cross-entropy information criterion (CIC): Cumulative version

$$\overline{\mathrm{CIC}}^{(t)}(d) = \bar{\mathcal{C}}^{(t-1)}\left(\widehat{\boldsymbol{\theta}}^{(t)}\right) + \widehat{\rho}\frac{d}{\sum_{s=0}^{t-1} n_s} \tag{18}$$

for $t = 1, \ldots, \tau$, where $\widehat{\boldsymbol{\theta}}^{(t)} := \underset{\boldsymbol{\theta}\in\boldsymbol{\Theta}}{\operatorname{argmin}}\, \bar{\mathcal{C}}^{(t-1)}(\boldsymbol{\theta})$ in (16). $\widehat{\rho}$ is a consistent estimator of ρ, such as the estimators in (19) and (20).

As t increases, the accumulated sample size $\sum_{s=0}^{t-1} n_s$ increases so that the free parameter dimension d can increase. Thus, the cumulative version of the CIC allows the use of a highly complex model if it can better approximate Q^*.

As a consistent and unbiased estimator of ρ, we can use

$$\widehat{\rho}^{(0)} = \frac{1}{n_0}\sum_{i=1}^{n_0} \frac{r\left(\mathbf{X}_i^{(0)}\right)}{q_{\widehat{\boldsymbol{\theta}}^{(0)}}\left(\mathbf{X}_i^{(0)}\right)} \tag{19}$$

at the 1^{st} iteration. At the t^{th} iteration for $t = 2, \ldots, \tau$, we can use

$$\widehat{\rho}^{(t-1)} = \frac{1}{\sum_{s=1}^{t-1} n_s}\sum_{s=1}^{t-1}\sum_{i=1}^{n_s} \frac{r\left(\mathbf{X}_i^{(s)}\right)}{q_{\widehat{\boldsymbol{\theta}}^{(s)}}\left(\mathbf{X}_i^{(s)}\right)}, \tag{20}$$

where we do not use the data at the 0^{th} iteration because they could potentially increase the variance of the resulting estimator if $Q_{\widehat{\boldsymbol{\theta}}^{(0)}} = Q_{\boldsymbol{\eta}}$ is too different from Q^*. The estimator in (20) takes the form of importance sampling estimator of $\rho = \mathbb{E}_\mu[r]$. The potential for the increased variance has been well studied in the importance sampling literature (e.g., Hesterberg, 1995; Owen and Zhou, 2000).

4.2 Mixture model and an EM algorithm

For a parametric family to have a flexible parameter dimension, we can consider a mixture model. Parametric mixture models are especially often used to approximate the optimal importance sampling density (Botev et al., 2013; Kurtz and Song, 2013; Wang and Zhou, 2015). Because the performance of importance sampling strongly depends on the approximate density, we need to carefully choose the number of mixture

components, k. We note that prior studies either assume that k is given (Botev et al., 2013; Kurtz and Song, 2013) or use a rule of thumb based on "some understanding of the structure of the problem at hand" (Wang and Zhou, 2015).

In theory, we can use the CIC to select the number of components, k, for any parametric mixture model, considering various parametric component families (e.g., exponential families are especially convenient because the MCE can be found by using the expectation-maximization (EM) algorithm (Dempster et al., 1977)). In this paper, we use the Gaussian mixture model (GMM) as the parametric family in which we find a density approximating the target density. Specifically, the GMM density can be expressed as

$$q(\mathbf{x};\boldsymbol{\theta}) = \sum_{j=1}^{k} \alpha_j \, q_j\big(\mathbf{x};\boldsymbol{\mu}_j,\boldsymbol{\Sigma}_j\big)\,, \tag{21}$$

where the component weights, $\alpha_j > 0, j = 1,\ldots,k$, satisfy $\sum_{j=1}^{k}\alpha_j = 1$. The jth Gaussian component density q_j is parametrized by the mean $\boldsymbol{\mu}_j$ and the covariance $\boldsymbol{\Sigma}_j$. Thus, the model parameter $\boldsymbol{\theta}$ denotes $(\alpha_1,\ldots,\alpha_k,\boldsymbol{\mu}_1,\ldots,\boldsymbol{\mu}_k,\boldsymbol{\Sigma}_1,\ldots,\boldsymbol{\Sigma}_k)$.

To find the MCE $\widehat{\boldsymbol{\theta}}^{(t)}$ of $\boldsymbol{\theta}$, we want to minimize $\bar{\mathcal{C}}^{(t-1)}(\boldsymbol{\theta})$ in (17) and thus set its gradient to zero:

$$-\frac{1}{\sum_{s=0}^{t-1} n_s}\sum_{s=0}^{t-1}\sum_{i=1}^{n_s}\nabla_{\boldsymbol{\theta}} h\Big(\mathbf{X}_i^{(s)},\widehat{\boldsymbol{\theta}}^{(s)},\boldsymbol{\theta}\Big) = 0.$$

This leads to the following updating equations for our version of the EM algorithm:

$$\alpha_j = \frac{\sum_{s=0}^{t-1}\sum_{i=1}^{n_s} \frac{r\left(\mathbf{X}_i^{(s)}\right)}{q_{\widehat{\boldsymbol{\theta}}^{(s)}}\left(\mathbf{X}_i^{(s)}\right)}\gamma_{ij}^{(s)}}{\sum_{s=0}^{t-1}\sum_{i=1}^{n_s} \frac{r\left(\mathbf{X}_i^{(s)}\right)}{q_{\widehat{\boldsymbol{\theta}}^{(s)}}\left(\mathbf{X}_i^{(s)}\right)}}, \tag{22}$$

$$\boldsymbol{\mu}_j = \frac{\sum_{s=0}^{t-1}\sum_{i=1}^{n_s} \frac{r\left(\mathbf{X}_i^{(s)}\right)}{q_{\widehat{\boldsymbol{\theta}}^{(s)}}\left(\mathbf{X}_i^{(s)}\right)}\gamma_{ij}^{(s)}\mathbf{X}_i^{(s)}}{\sum_{s=0}^{t-1}\sum_{i=1}^{n_s} \frac{r\left(\mathbf{X}_i^{(s)}\right)}{q_{\widehat{\boldsymbol{\theta}}^{(s)}}\left(\mathbf{X}_i^{(s)}\right)}\gamma_{ij}^{(s)}}, \tag{23}$$

$$\boldsymbol{\Sigma}_j = \frac{\sum_{s=0}^{t-1}\sum_{i=1}^{n_s} \frac{r\left(\mathbf{X}_i^{(s)}\right)}{q_{\widehat{\boldsymbol{\theta}}^{(s)}}\left(\mathbf{X}_i^{(s)}\right)}\gamma_{ij}^{(s)}(\mathbf{X}_i^{(s)}-\boldsymbol{\mu}_j)(\mathbf{X}_i^{(s)}-\boldsymbol{\mu}_j)^T}{\sum_{s=0}^{t-1}\sum_{i=1}^{n_s} \frac{r\left(\mathbf{X}_i^{(s)}\right)}{q_{\widehat{\boldsymbol{\theta}}^{(s)}}\left(\mathbf{X}_i^{(s)}\right)}\gamma_{ij}^{(s)}}, \tag{24}$$

where

$$\gamma_{ij}^{(s)} = \frac{\alpha_j\, q_j\left(\mathbf{X}_i^{(s)};\boldsymbol{\mu}_j,\boldsymbol{\Sigma}_j\right)}{\sum_{j'=1}^{k}\alpha_{j'}\, q_{j'}\left(\mathbf{X}_i^{(s)};\boldsymbol{\mu}_{j'},\boldsymbol{\Sigma}_{j'}\right)}. \tag{25}$$

The right-hand sides of the updating equations in (22), (23), and (24) involve $\boldsymbol{\theta} = (\alpha_1,\ldots,\alpha_k,\boldsymbol{\mu}_1,\ldots,\boldsymbol{\mu}_k,\boldsymbol{\Sigma}_1,\ldots,\boldsymbol{\Sigma}_k)$ either explicitly or implicitly through $\gamma_{ij}^{(s)}$. The equations cannot be analytically solved for $\boldsymbol{\theta}$. Instead, starting with an initial guess of $\boldsymbol{\theta}$, our version of the EM algorithm alternates between the expectation step (computing $\gamma_{ij}^{(s)}$) and the maximization step (updating $\boldsymbol{\theta}$) based on the updating equations until convergence is reached.

We note that Kurtz and Song (2013) derive similar, but simpler, updating equations for the purpose of minimizing the cross-entropy, although they do not use the equations for the EM algorithm. Prior studies (Botev et al., 2013; Wang and Zhou, 2015; Kurtz and Song, 2013) using mixture models for the cross-entropy method for importance sampling (recall that this method is a special case of the procedure in Figure 1 when r is proportional to the optimal importance sampling density) do not iterate their updating equations; rather, they solve them only once when new data are gathered in each iteration. This paper uses the aforementioned EM algorithm (i.e., iterating the updating equations until convergence) within each iteration to minimize $\bar{\mathcal{C}}^{(t-1)}(\boldsymbol{\theta})$ in (17) (see Appendix B for the implementation details).

4.3 Summary of the CIC-based distribution approximation procedure

Similar to the AIC, the CIC tends to decrease and then slowly increase as d increases, subject to the randomness of the data (see Figure 2, where k is the number of mixture components in the GMM with unconstrained means and covariances. k is proportional to the free parameter dimension $d = (k-1) + k(p + p(p+1)/2)$, with p denoting the dimension of the GMM density support). Thus, we minimize the CIC to find the best number of components, $k^{*(t)}$ (or the best model dimension $d^{*(t)}$) to use in the t^{th} iteration (see Appendix B for the implementation details). Figure 3 summarizes the CIC-based distribution approximation procedure.

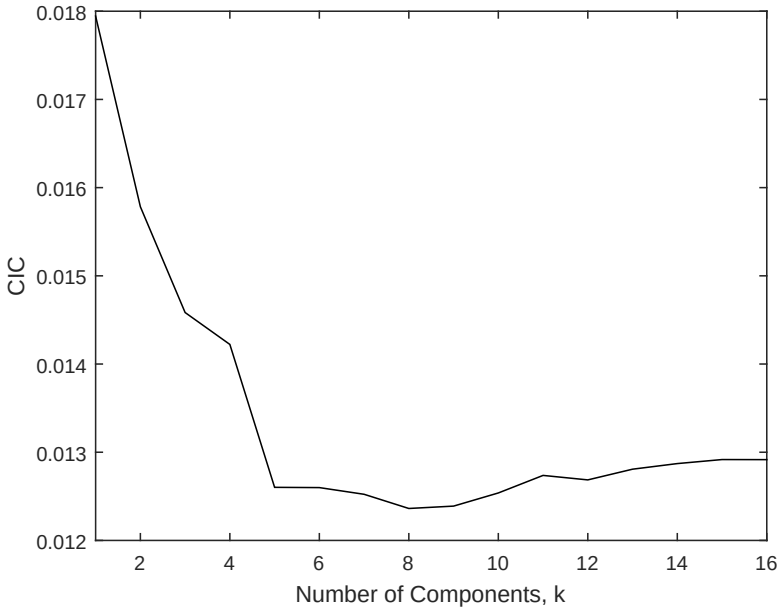

Figure 2: CIC observed in the example in Section 5

If we additionally want to estimate a quantity of interest such as ρ, we can sample $\mathbf{X}_1^{(\tau)}, \ldots, \mathbf{X}_n^{(\tau)} \sim Q_{\widehat{\boldsymbol{\theta}}^{(\tau)}}$ and use the estimator such as $\widehat{\rho}^{(\tau)}$ in (20). Section 5 uses this additional step for importance sampling.

CIC-Based Approximation of the Target Distribution Q^*

Inputs: iteration counter $t = 1$, the maximum iteration τ, and the initial parameter $\widehat{\boldsymbol{\theta}}^{(0)} = \boldsymbol{\eta} \in \boldsymbol{\Theta}$.

1. Sample $\mathbf{X}_1^{(t-1)}, \ldots, \mathbf{X}_n^{(t-1)} \sim Q_{\widehat{\boldsymbol{\theta}}^{(t-1)}}$.
2. Find the best model dimension $d^{*(t)} := \underset{d \geq 1}{\operatorname{argmin}} \overline{\mathrm{CIC}}^{(t)}(d)$ to use, where $\overline{\mathrm{CIC}}^{(t)}(d) := \bar{\mathcal{C}}^{(t-1)}\left(\widehat{\boldsymbol{\theta}}^{(t)}\right) + \widehat{\rho}\frac{d}{\sum_{s=0}^{t-1} n_s}$ is in (18) with the MCE $\widehat{\boldsymbol{\theta}}^{(t)} := \underset{\boldsymbol{\theta} \in \boldsymbol{\Theta}}{\operatorname{argmin}} \bar{\mathcal{C}}^{(t-1)}(\boldsymbol{\theta})$ in (16). $\widehat{\rho}$ is a consistent estimator of ρ, such as the estimators in (19) and (20).
3. If $t = \tau$, output the approximate distribution $Q_{\widehat{\boldsymbol{\theta}}^{(\tau)}}$. Otherwise, increment t by 1 and go to Step 1.

Figure 3: CIC-based approximation of the target distribution Q^*

5 Numerical Study

Theoretically, importance sampling (Kahn and Marshall, 1953), which has been widely used to reduce the estimator variance in Monte Carlo studies (Givens and Raftery, 1996; Zhang, 1996; Owen and Zhou, 2000; Neddermeyer, 2009), can lead to zero variance of an estimator if we can sample from the optimal distribution. In practice, however, we do not know the optimal distribution, but usually we can evaluate a nonnegative function proportional to the optimal density. Therefore, we use the CIC to approximate the optimal distribution. We expect that if the approximation is better, the variance reduction will be greater.

As a benchmark, we consider the importance sampling method in Kurtz and Song (2013), which has a similar structure to our CIC-based approach. Their method, which is called cross-entropy-based adaptive importance sampling using Gaussian mixture (CE-AIS-GM) uses the GMM with a pre-specified value for the number of mixture components, k. The GMM parameters are updated once in each iteration using updating equations similar to ours in (22), (23), and (24) to reduce the cross-entropy from the GMM density to the optimal importance sampling density.

We use a classical example in the structural safety literature, which is also used in Kurtz and Song (2013). For $\mathbf{X}$ with the density $\phi(\mathbf{x})$ of a bivariate Gaussian

distribution with zero mean and identity covariance matrix, a system of interest fails when $\mathbf{X}$ falls on the region $\{\mathbf{x} \in \mathbb{R}^2 : g(\mathbf{x}) \leq 0\}$, where

$$g(\mathbf{x}) = b - x_2 - \kappa \left(x_1 - e\right)^2 . \tag{26}$$

To compare the CE-AIS-GM with our method, which we call CIC-IS, we vary the parameter $b = 1.5$, 2.0, and 2.5, to test three different failure thresholds. We fix the other two parameters, $\kappa = 0.1$ and $e = 0$, to maintain the shape of the failure region. We use the same sample size in Kurtz and Song (2013), namely, the total of 8700 replications: $n_t = 1000$ for $t = 0, \ldots, 6$ and $n_\tau = 1700$ for $\tau = 7$.

As is well known in the importance sampling literature (Kahn and Marshall, 1953), the optimal importance sampling density $q^*(\mathbf{x})$ for this example should be proportional to $r(\mathbf{x}) = \phi(\mathbf{x})\mathbb{I}(g(\mathbf{x}) \leq 0)$, where $\mathbb{I}(\cdot)$ is the indicator function. Therefore, $\rho = \int r \, \mathrm{d}\mu$ is the probability of the failure event.

We set the CE-AIS-GM to use $k = 30$ and to estimate the failure probability based only on the last ($\tau^{(th)}$) iteration data as in Kurtz and Song (2013). The CIC-IS adaptively chooses k within the algorithm described in Figure 3 and uses the data from all iterations except the $0^{(th)}$ iteration to estimate the failure probability by $\widehat{\rho}^{(\tau)}$ in (20), as described in Section 4. Because the CIC helps find a distribution reasonably close to the optimal distribution throughout all iterations, in general we can use the data from all iterations, not just the last iteration, to estimate the failure probability. It also helps to reduce the standard error of the estimator.

Table 1 shows the estimation results based on 500 experiment repetitions. The sample mean of the failure probability estimates (Mean) decreases as the threshold, b, increases. Compared to the CE-AIS-GM, the CIC-IS obtains at least 50% smaller standard errors.

Figure 4 compares the theoretically optimal density q^* and a CIC-IS density $q_{\widehat{\boldsymbol{\theta}}^{(\tau)}}$, for $b = 1.5$. We observe that the CIC-IS density with the automatically chosen $k = 10$ is close to the theoretically optimal density, being able to capture the shape of the important region.

Table 1: Comparison between CE-AIS-GM and CIC-IS

b	Method	Mean	Standard Error
1.5	CE-AIS-GM	0.082902	0.001145
	CIC-IS	0.082911	0.000506
2.0	CE-AIS-GM	0.030174	0.000526
	CIC-IS	0.030173	0.000213
2.5	CE-AIS-GM	0.008908	0.000211
	CIC-IS	0.008910	0.000099

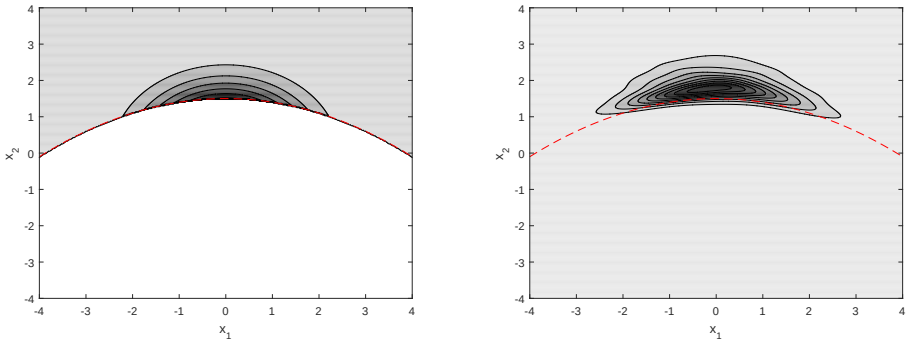

(a) Optimal density (b) CIC-IS density ($k = 10$)

Figure 4: Comparison between the theoretically optimal density q^* and the CIC-IS density $q_{\widehat{\boldsymbol{\theta}}^{(\tau)}}$. The red dashed line is the failure boundary $g(\mathbf{x}) = 0$ when $b = 1.5$.

6 Summary and Future Research Directions

This paper proposed an information criterion, called CIC, to find a parametric density that has the asymptotically minimum cross-entropy to a target density to approximate. The CIC is the sum of two terms: an estimator of the cross-entropy (up to a multiplicative constant) from the parametric density to the target density, and a model complexity penalty term that is proportional to the free parameter dimension of the parametric density. Under certain regularity conditions, we proved that the penalty term corrected the asymptotic bias of the first term in estimating the true cross-entropy. Empirically, we demonstrated that minimizing the CIC could lead to a density that well approximated the optimal importance sampling density.

Our findings suggest several future research directions. Importance sampling has been used (Liu, 1996; Bassetti and Diaconis, 2006) as an alternative to the Metropolis-Hastings algorithm (Metropolis et al., 1953; Hastings, 1970) for Bayesian inference. A CIC-based importance sampling may lead to more accurate inference with less

computational burden than existing importance sampling methods using a fixed trial (or proposal) density.

Variational inference (Jordan et al., 1999; Wainwright and Jordan, 2008), which is an emerging Bayesian inference technique for approximating a posterior distribution, is a fast and scalable alternative to Markov chain Monte Carlo (MCMC) (Hastings, 1970; Gelfand and Smith, 1990) (for a recent review on variational inference and its comparison to MCMC, see Blei et al. (2016)). The computational efficiency is achieved by positing the posterior approximation as an optimization problem that minimizes the KL divergence from a distribution in a parametric family to the posterior distribution. The computational complexity of the optimization problem depends on the complexity of the parametric family (Blei et al., 2016). We conjecture that the CIC could be a way to balance the inference accuracy and the computational complexity in variational inference.

In proposing the AIC, Akaike (1974) stated that "IC stands for information criterion and A is added so that similar statistics, BIC, DIC etc., may follow." Indeed, the BIC (Schwarz, 1978) soon followed. Now the CIC, which can be viewed as a generalization of the AIC, follows. Future research on a Bayesian criterion that is to the CIC as the BIC is to the AIC appears to be next, with an obvious choice of the criterion name.

Appendix A: Assumptions and Proofs

Having stated Assumption 1 in Section 3.3, we give the remaining assumptions that constitute the regularity conditions for our theoretical results. Specifically, Assumptions 2–8 are the standard regularity conditions to establish consistency (see Lemma 1) and asymptotic normality (see Lemma 2) of an M-estimator. Assumptions 9 and 10 are mild conditions to regularize $q_{\boldsymbol{\theta}}$ in a neighborhood of $\boldsymbol{\theta}^*$ to establish Corollary 4.

Theorem 5 is based on the additional regularity conditions, Assumptions 11–16. Assumption 11 is a common regularity condition to interchange expectation and differentiation. Assumptions 12 and 15 are mild conditions for some matrices involving Hessians to converge in probability. Assumptions 13, 14, and 16 regulate the

MCE $\widehat{\boldsymbol{\theta}}_n^{(t)}$ in (10) for $n \geq 1$ and $t = 1, \ldots, \tau$. The uniform integrability conditions in Assumptions 13 and 16 are commonly used (e.g., Conditions A7–A8 in Donohue et al. (2011)) for an information criterion similar to the AIC to make the model complexity penalty to be expressed in the free parameter dimension d. Assumption 14 is relatively mild, because $\widehat{\boldsymbol{\theta}}_n^{(t)}$ approaches $\boldsymbol{\theta}^*$, which is an interior point under Assumption 4.

We recall that the functions, $\mathcal{C}(\boldsymbol{\theta}) := -\mathbb{E}_\mu[r \log q_{\boldsymbol{\theta}}]$ and $h(\mathbf{X}, \boldsymbol{\eta}, \boldsymbol{\theta}) := \frac{r(\mathbf{X})}{q_{\boldsymbol{\eta}}(\mathbf{X})} \log q_{\boldsymbol{\theta}}(\mathbf{X})$, are defined in (6) and (8), respectively. Below, $\nabla_{\boldsymbol{\theta}}$ and $\nabla^2_{\boldsymbol{\theta}}$ denote the gradient and the Hessian with respect to $\boldsymbol{\theta}$, respectively. For example, $\nabla_{\boldsymbol{\theta}} h(\mathbf{X}, \boldsymbol{\eta}, \boldsymbol{\theta}^*)$ denotes the gradient of $h(\mathbf{X}, \boldsymbol{\eta}, \boldsymbol{\theta})$ with respect to $\boldsymbol{\theta}$ at $\boldsymbol{\theta}^*$.

Assumption 2. *For any $\boldsymbol{\eta} \in \boldsymbol{\Theta}$, $h(\mathbf{x}, \boldsymbol{\eta}, \cdot)$ is continuous on $\boldsymbol{\Theta}$ for a.e. $\mathbf{x}$ under $Q_{\boldsymbol{\eta}}$.*

Assumption 3. *For any $\boldsymbol{\eta} \in \boldsymbol{\Theta}$, there exists a measurable function $g_{\boldsymbol{\eta}}$ such that $\int |g_{\boldsymbol{\eta}}| \, \mathrm{d}Q_{\boldsymbol{\eta}} < \infty$ and $|h(\mathbf{x}, \boldsymbol{\eta}, \boldsymbol{\theta})| \leq g_{\boldsymbol{\eta}}(\mathbf{x})$ for a.e. $\mathbf{x}$ under $Q_{\boldsymbol{\eta}}$ and all $\boldsymbol{\theta} \in \boldsymbol{\Theta}$.*

Assumption 4. *$\boldsymbol{\theta}^* := \operatorname{argmin}_{\boldsymbol{\theta} \in \boldsymbol{\Theta}} \mathcal{C}(\boldsymbol{\theta})$ is an interior point of $\boldsymbol{\Theta}$.*

Assumption 5. *For any $\boldsymbol{\eta} \in \boldsymbol{\Theta}$, $h(\mathbf{x}, \boldsymbol{\eta}, \cdot)$ is twice continuously differentiable in a neighborhood of $\boldsymbol{\theta}^*$ for a.e. $\mathbf{x}$ under $Q_{\boldsymbol{\eta}}$.*

Assumption 6. *For any $\boldsymbol{\eta} \in \boldsymbol{\Theta}$, there exist measurable functions $g^{(i)}_{\boldsymbol{\eta}}, i = 1, 2$ such that $\int |g^{(i)}_{\boldsymbol{\eta}}| \, \mathrm{d}Q_{\boldsymbol{\eta}} < \infty, i = 1, 2$ and $\|\nabla^i_{\boldsymbol{\theta}} h(\mathbf{x}, \boldsymbol{\eta}, \boldsymbol{\theta})\| \leq g^{(i)}_{\boldsymbol{\eta}}(\mathbf{x}), i = 1, 2$ for a.e. $\mathbf{x}$ under $Q_{\boldsymbol{\eta}}$ and all $\boldsymbol{\theta} \in \boldsymbol{\Theta}$. The norm is Euclidean for $i = 1$ and Frobenius for $i = 2$.*

Assumption 7. *$\boldsymbol{\Gamma} := -\mathbb{E}_\mu[r \nabla^2_{\boldsymbol{\theta}} \log q_{\boldsymbol{\theta}^*}]$ is nonsingular.*

Assumption 8. *For any $\boldsymbol{\eta} \in \boldsymbol{\Theta}$, $\boldsymbol{\Lambda}_{\boldsymbol{\eta}} := \mathbb{E}_{Q_{\boldsymbol{\eta}}}\left[\nabla_{\boldsymbol{\theta}} h(\mathbf{X}, \boldsymbol{\eta}, \boldsymbol{\theta}^*) \nabla_{\boldsymbol{\theta}} h(\mathbf{X}, \boldsymbol{\eta}, \boldsymbol{\theta}^*)^T\right]$ exists.*

Assumption 9. *$\nabla_{\boldsymbol{\theta}}(1/q_{\boldsymbol{\theta}}(\mathbf{x}))$ is continuous in a neighborhood of $\boldsymbol{\theta}^*$ for a.e. $\mathbf{x}$ under μ.*

Assumption 10. *$\left\| \mathbb{E}_\mu\left[\nabla_{\boldsymbol{\theta}}(1/q_{\boldsymbol{\theta}}) \, \nabla_{\boldsymbol{\theta}} q_{\boldsymbol{\theta}^*} \, (\nabla_{\boldsymbol{\theta}} q_{\boldsymbol{\theta}^*})^T\right] \right\|$ is bounded in a neighborhood of $\boldsymbol{\theta}^*$.*

Assumption 11. *There exists a measurable function $g_0(\mathbf{x})$ such that $\int |g_0| \, \mathrm{d}\mu < \infty$ and $\|\nabla_{\boldsymbol{\theta}} q_{\boldsymbol{\theta}}(\mathbf{x})\| \leq g_0(\mathbf{x})$ in a neighborhood of $\boldsymbol{\theta}^*$ for a.e. $\mathbf{x}$ under μ.*

Assumption 12. *$-\mathbb{E}_\mu\left[r \nabla^2_{\boldsymbol{\theta}} \log q_{\boldsymbol{\theta}}\right]$ is continuous at $\boldsymbol{\theta}^*$.*

Assumption 13. $n\left(\widehat{\boldsymbol{\theta}}_n^{(t)}-\boldsymbol{\theta}^*\right)^T\left(-\mathbb{E}_\mu\left[r\nabla^2_{\boldsymbol{\theta}}\log q_{\boldsymbol{\theta}}\right]\right)\left(\widehat{\boldsymbol{\theta}}_n^{(t)}-\boldsymbol{\theta}^*\right)$, $n\geq 1$, *are uniformly integrable for* $t=2,\ldots,\tau$ *in a neighborhood of* $\boldsymbol{\theta}^*$.

Assumption 14. $\widehat{\boldsymbol{\theta}}_n^{(t)}:=\underset{\boldsymbol{\theta}\in\boldsymbol{\Theta}}{\operatorname{argmin}}\,\bar{\mathcal{C}}_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}(\boldsymbol{\theta})$ *is an interior point of* $\boldsymbol{\Theta}$ *for* $n\geq 1$ *and* $t=1,\ldots,\tau$.

Assumption 15. *For any* $\boldsymbol{\eta}\in\boldsymbol{\Theta}$, $\|\nabla_{\boldsymbol{\eta}}[\nabla^2_{\boldsymbol{\theta}}h(\mathbf{x},\boldsymbol{\eta},\boldsymbol{\theta})]_{kl}\|$ *and* $\|\nabla_{\boldsymbol{\theta}}[\nabla^2_{\boldsymbol{\theta}}h(\mathbf{x},\boldsymbol{\eta},\boldsymbol{\theta})]_{kl}\|$ *are bounded in a neighborhood of* $\boldsymbol{\theta}^*$ *for a.e.* $\mathbf{x}$ *under* $Q_{\boldsymbol{\eta}}$ *for* $k=1,\ldots,d$ *and* $l=1,\ldots,d$.

Assumption 16. $\left(\widehat{\boldsymbol{\theta}}_n^{(t)}-\boldsymbol{\theta}^*\right)^T\left(-\sum_{i=1}^n\nabla^2_{\boldsymbol{\theta}}h\left(\mathbf{X}_i^{(t-1)},\widehat{\boldsymbol{\theta}}_n^{(t-1)},\boldsymbol{\theta}\right)\right)\left(\widehat{\boldsymbol{\theta}}_n^{(t)}-\boldsymbol{\theta}^*\right)$, $n\geq 1$, *are uniformly integrable for* $t=2,\ldots,\tau$ *in a neighborhood of* $\boldsymbol{\theta}^*$.

Below, we provide the proofs of Corollary 4 and Theorem 5.

Proof of Corollary 4. We first express $\boldsymbol{\Gamma}$ in an equivalent form under the given conditions and show that $\boldsymbol{\Lambda}_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}$ approaches $\rho\boldsymbol{\Gamma}$ in probability as $n\to\infty$. We have

$$\begin{aligned}
\boldsymbol{\Gamma}&=-\mathbb{E}_\mu\left[r\nabla^2_{\boldsymbol{\theta}}\log q_{\boldsymbol{\theta}^*}\right]\\
&=-\mathbb{E}_\mu\left[r\left(-\frac{1}{q^2_{\boldsymbol{\theta}^*}}\nabla_{\boldsymbol{\theta}}q_{\boldsymbol{\theta}^*}\left(\nabla_{\boldsymbol{\theta}}q_{\boldsymbol{\theta}^*}\right)^T+\frac{1}{q_{\boldsymbol{\theta}^*}}\nabla^2_{\boldsymbol{\theta}}q_{\boldsymbol{\theta}^*}\right)\right]\\
&=-\rho\mathbb{E}_\mu\left[-\frac{1}{q_{\boldsymbol{\theta}^*}}\nabla_{\boldsymbol{\theta}}q_{\boldsymbol{\theta}^*}\left(\nabla_{\boldsymbol{\theta}}q_{\boldsymbol{\theta}^*}\right)^T+\nabla^2_{\boldsymbol{\theta}}q_{\boldsymbol{\theta}^*}\right] \qquad (27)\\
&=\rho\mathbb{E}_\mu\left[\frac{1}{q_{\boldsymbol{\theta}^*}}\nabla_{\boldsymbol{\theta}}q_{\boldsymbol{\theta}^*}\left(\nabla_{\boldsymbol{\theta}}q_{\boldsymbol{\theta}^*}\right)^T\right], \qquad (28)
\end{aligned}$$

where the equation in (27) holds, because $q_{\boldsymbol{\theta}^*}=q^*=r/\rho$ by Assumption 1. The equation in (28) is because $\mathbb{E}_\mu\nabla^2_{\boldsymbol{\theta}}q_{\boldsymbol{\theta}^*}=\nabla^2_{\boldsymbol{\theta}}\mathbb{E}_\mu q_{\boldsymbol{\theta}^*}=0$ by the dominated convergence theorem under Assumption 6.

On the other hand,

$$\begin{aligned}
\boldsymbol{\Lambda}_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}} &= \mathbb{E}_{Q_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}}\left[\nabla_{\boldsymbol{\theta}} h\left(\mathbf{X}, \widehat{\boldsymbol{\theta}}_n^{(t-1)}, \boldsymbol{\theta}^*\right) \nabla_{\boldsymbol{\theta}} h\left(\mathbf{X}, \widehat{\boldsymbol{\theta}}_n^{(t-1)}, \boldsymbol{\theta}^*\right)^T\right] \\
&= \mathbb{E}_{Q_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}}\left[\frac{r^2(\mathbf{X})}{q^2_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}(\mathbf{X})} \nabla_{\boldsymbol{\theta}} \log q_{\boldsymbol{\theta}^*}(\mathbf{X}) \left(\nabla_{\boldsymbol{\theta}} \log q_{\boldsymbol{\theta}^*}(\mathbf{X})\right)^T\right] \\
&= \mathbb{E}_{\mu}\left[\frac{r^2}{q_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}} \nabla_{\boldsymbol{\theta}} \log q_{\boldsymbol{\theta}^*} \left(\nabla_{\boldsymbol{\theta}} \log q_{\boldsymbol{\theta}^*}\right)^T\right] \\
&= \mathbb{E}_{\mu}\left[\frac{r^2}{q_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}} q^2_{\boldsymbol{\theta}^*}} \nabla_{\boldsymbol{\theta}} q_{\boldsymbol{\theta}^*} \left(\nabla_{\boldsymbol{\theta}} q_{\boldsymbol{\theta}^*}\right)^T\right] \\
&= \rho^2 \mathbb{E}_{\mu}\left[\frac{1}{q_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}} \nabla_{\boldsymbol{\theta}} q_{\boldsymbol{\theta}^*} \left(\nabla_{\boldsymbol{\theta}} q_{\boldsymbol{\theta}^*}\right)^T\right] \qquad (29) \\
&= \rho^2 \int \left(\frac{1}{q_{\boldsymbol{\theta}^*}(\mathbf{x})} + \left(\widehat{\boldsymbol{\theta}}_n^{(t-1)} - \boldsymbol{\theta}^*\right)^T \nabla_{\boldsymbol{\theta}} \left(\frac{1}{q_{\tilde{\boldsymbol{\theta}}_n(\mathbf{x})}(\mathbf{x})}\right)\right) \nabla_{\boldsymbol{\theta}} q_{\boldsymbol{\theta}^*}(\mathbf{x}) \left(\nabla_{\boldsymbol{\theta}} q_{\boldsymbol{\theta}^*}(\mathbf{x})\right)^T \mathrm{d}\mu(\mathbf{x}) \qquad (30) \\
&= \rho \boldsymbol{\Gamma} + \rho^2 \int \left(\widehat{\boldsymbol{\theta}}_n^{(t-1)} - \boldsymbol{\theta}^*\right)^T \nabla_{\boldsymbol{\theta}} \left(\frac{1}{q_{\tilde{\boldsymbol{\theta}}_n(\mathbf{x})}(\mathbf{x})}\right) \nabla_{\boldsymbol{\theta}} q_{\boldsymbol{\theta}^*}(\mathbf{x}) \left(\nabla_{\boldsymbol{\theta}} q_{\boldsymbol{\theta}^*}(\mathbf{x})\right)^T \mathrm{d}\mu(\mathbf{x}) \\
&= \rho \boldsymbol{\Gamma} + o_p(1), \qquad (31)
\end{aligned}$$

where the equation in (29) holds by Assumption 1. $\tilde{\boldsymbol{\theta}}_n(\mathbf{x})$ is an intermediate point between $\widehat{\boldsymbol{\theta}}_n^{(t-1)}$ and $\boldsymbol{\theta}^*$ (i.e., $\left\|\tilde{\boldsymbol{\theta}}_n(\mathbf{x}) - \boldsymbol{\theta}^*\right\| \leq \left\|\widehat{\boldsymbol{\theta}}_n^{(t-1)} - \boldsymbol{\theta}^*\right\|$) such that the equality in (30) holds by Taylor expansion under Assumption 9. Because $\widehat{\boldsymbol{\theta}}_n^{(t-1)}$ converges in probability to $\boldsymbol{\theta}^*$ for $t = 2, \ldots, \tau$, the equation in (31) holds under Assumption 10.

By Lemma 2 and Slutsky's theorem, $\sqrt{n}\left(\widehat{\boldsymbol{\theta}}_n^{(t)} - \boldsymbol{\theta}^*\right)$ converges in distribution to $N\left(0, \rho\boldsymbol{\Gamma}^{-1}\right)$ as $n \to \infty$ for $t = 2, \ldots, \tau$.

□

Proof of Theorem 5.

We first simplify $\mathcal{C}\left(\widehat{\boldsymbol{\theta}}_n^{(t)}\right)$ and $\bar{\mathcal{C}}_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}\left(\widehat{\boldsymbol{\theta}}_n^{(t)}\right)$. We then use the simplified expressions to derive the bias of interest.

$$
\begin{aligned}
\mathcal{C}\left(\widehat{\boldsymbol{\theta}}_n^{(t)}\right) &= -\int r \log q_{\widehat{\boldsymbol{\theta}}_n^{(t)}} \, \mathrm{d}\mu \\
&= -\int \frac{r}{q_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}} q_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}} \log q_{\widehat{\boldsymbol{\theta}}_n^{(t)}} \, \mathrm{d}\mu \qquad (32) \\
&= -\mathbb{E}_{Q_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}}\left[h\left(\mathbf{X}, \widehat{\boldsymbol{\theta}}_n^{(t-1)}, \widehat{\boldsymbol{\theta}}_n^{(t)}\right)\right],
\end{aligned}
$$

where $\mathbf{X} \sim Q_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}$. The equation in (32) holds under the condition $Q^* \ll Q_{\boldsymbol{\theta}}$ for all $\boldsymbol{\theta} \in \boldsymbol{\Theta}$, because $q_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}(\mathbf{x}) = 0$ implies $r(\mathbf{x}) = 0$ for any $\mathbf{x}$.

We take a second-order expansion of $\mathcal{C}\left(\widehat{\boldsymbol{\theta}}_n^{(t)}\right)$ about $\boldsymbol{\theta}^*$:

$$
\begin{aligned}
\mathcal{C}\left(\widehat{\boldsymbol{\theta}}_n^{(t)}\right) = -\mathbb{E}_{Q_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}}\left[h\left(\mathbf{X}, \widehat{\boldsymbol{\theta}}_n^{(t-1)}, \boldsymbol{\theta}^*\right)\right] - \left(\widehat{\boldsymbol{\theta}}_n^{(t)} - \boldsymbol{\theta}^*\right)^T \mathbb{E}_{Q_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}}\left[\nabla_{\boldsymbol{\theta}} h\left(\mathbf{X}, \widehat{\boldsymbol{\theta}}_n^{(t-1)}, \boldsymbol{\theta}^*\right)\right] \\
- \frac{1}{2}\left(\widehat{\boldsymbol{\theta}}_n^{(t)} - \boldsymbol{\theta}^*\right)^T \mathbb{E}_{Q_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}}\left[\nabla^2_{\boldsymbol{\theta}} h\left(\mathbf{X}, \widehat{\boldsymbol{\theta}}_n^{(t-1)}, \tilde{\boldsymbol{\theta}}_n(\mathbf{X})\right)\right]\left(\widehat{\boldsymbol{\theta}}_n^{(t)} - \boldsymbol{\theta}^*\right), \qquad (33)
\end{aligned}
$$

where $\tilde{\boldsymbol{\theta}}_n(\mathbf{X})$ is an intermediate point between $\widehat{\boldsymbol{\theta}}_n^{(t)}$ and $\boldsymbol{\theta}^*$ (i.e., $\left\|\tilde{\boldsymbol{\theta}}_n(\mathbf{X}) - \boldsymbol{\theta}^*\right\| \leq \left\|\widehat{\boldsymbol{\theta}}_n^{(t)} - \boldsymbol{\theta}^*\right\|$) and passing the derivative inside the expectation is valid under Assumptions 5 and 6. In (33), the zeroth-order term is $\mathcal{C}(\boldsymbol{\theta}^*)$ by definition and the first-order term is zero because

$$
\begin{aligned}
\mathbb{E}_{Q_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}}\left[\nabla_{\boldsymbol{\theta}} h\left(\mathbf{X}, \widehat{\boldsymbol{\theta}}_n^{(t-1)}, \boldsymbol{\theta}^*\right)\right] &= \mathbb{E}_{Q_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}}\left[\frac{r(\mathbf{X})}{q_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}(\mathbf{X})} \nabla_{\boldsymbol{\theta}} \log q_{\boldsymbol{\theta}^*}(\mathbf{X})\right] \\
&= \mathbb{E}_{\mu}\left[\frac{r(\mathbf{X})}{q_{\boldsymbol{\theta}^*}(\mathbf{X})} \nabla_{\boldsymbol{\theta}} q_{\boldsymbol{\theta}^*}(\mathbf{X})\right] \\
&= \rho \mathbb{E}_{\mu}\left[\nabla_{\boldsymbol{\theta}} q_{\boldsymbol{\theta}^*}(\mathbf{X})\right] \qquad (34) \\
&= \rho \nabla_{\boldsymbol{\theta}} \mathbb{E}_{\mu}\left[q_{\boldsymbol{\theta}^*}(\mathbf{X})\right] \qquad (35) \\
&= 0,
\end{aligned}
$$

where the equation in (34) holds under Assumption 1 and the interchange of expectation and differentiation in (35) holds under Assumption 11. Therefore, by defining $\boldsymbol{\delta}_n := \sqrt{n}\left(\widehat{\boldsymbol{\theta}}_n^{(t)} - \boldsymbol{\theta}^*\right)$ and $\tilde{\boldsymbol{\Gamma}}_n := -\mathbb{E}_{\mu}\left[r \nabla^2_{\boldsymbol{\theta}} \log q_{\tilde{\boldsymbol{\theta}}_n}\right]$, the expression of $\mathcal{C}\left(\widehat{\boldsymbol{\theta}}_n^{(t)}\right)$ in (33)

is simplified to

$$\begin{aligned}\mathcal{C}\left(\widehat{\boldsymbol{\theta}}_n^{(t)}\right) &= \mathcal{C}(\boldsymbol{\theta}^*) - \frac{1}{2n}\boldsymbol{\delta}_n^T \mathbb{E}_{Q_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}}\left[\nabla^2_{\boldsymbol{\theta}} h\left(\mathbf{X}, \widehat{\boldsymbol{\theta}}_n^{(t-1)}, \tilde{\boldsymbol{\theta}}_n(\mathbf{X})\right)\right]\boldsymbol{\delta}_n \\ &= \mathcal{C}(\boldsymbol{\theta}^*) + \frac{1}{2n}\boldsymbol{\delta}_n^T \tilde{\boldsymbol{\Gamma}}_n \boldsymbol{\delta}_n.\end{aligned}$$

In the second term, $\tilde{\boldsymbol{\Gamma}}_n$ converges in probability to $\boldsymbol{\Gamma}$ as $n \to \infty$ by the continuous mapping theorem under Assumption 12, because the intermediate point $\tilde{\boldsymbol{\theta}}_n(\mathbf{X})$ between $\widehat{\boldsymbol{\theta}}_n^{(t)}$ and $\boldsymbol{\theta}^*$ converges in probability to $\boldsymbol{\theta}^*$ under μ.

Also, because $\boldsymbol{\delta}_n$ converges in distribution to $N\left(0, \rho\boldsymbol{\Gamma}^{-1}\right)$, Slutsky's theorem yields that $\boldsymbol{\delta}_n^T \tilde{\boldsymbol{\Gamma}}_n \boldsymbol{\delta}_n$ converges in distribution to $\rho\chi^2_d$ as $n \to \infty$. Under Assumption 13 that imposes uniform integrability on $\boldsymbol{\delta}_n^T \tilde{\boldsymbol{\Gamma}}_n \boldsymbol{\delta}_n$, it follows that

$$\begin{aligned}\mathbb{E}\left[\mathcal{C}\left(\widehat{\boldsymbol{\theta}}_n^{(t)}\right)\right] &= \mathcal{C}(\boldsymbol{\theta}^*) + \frac{1}{2n}\mathbb{E}\left[\boldsymbol{\delta}_n^T \tilde{\boldsymbol{\Gamma}}_n \boldsymbol{\delta}_n\right] \\ &= \mathcal{C}(\boldsymbol{\theta}^*) + \frac{\rho d}{2n} + o\left(\frac{1}{n}\right). \qquad (36)\end{aligned}$$

For similar simplification, we take a second-order expansion of $\bar{\mathcal{C}}_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}\left(\widehat{\boldsymbol{\theta}}_n^{(t)}\right)$ about $\boldsymbol{\theta}^*$:

$$\begin{aligned}\bar{\mathcal{C}}_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}\left(\widehat{\boldsymbol{\theta}}_n^{(t)}\right) &= -\frac{1}{n}\sum_{i=1}^n h\left(\mathbf{X}_i^{(t-1)}, \widehat{\boldsymbol{\theta}}_n^{(t-1)}, \widehat{\boldsymbol{\theta}}_n^{(t)}\right) \\ &= -\frac{1}{n}\sum_{i=1}^n \left(h\left(\mathbf{X}_i^{(t-1)}, \widehat{\boldsymbol{\theta}}_n^{(t-1)}, \boldsymbol{\theta}^*\right) + \left(\widehat{\boldsymbol{\theta}}_n^{(t)} - \boldsymbol{\theta}^*\right)^T \nabla_{\boldsymbol{\theta}} h\left(\mathbf{X}_i^{(t-1)}, \widehat{\boldsymbol{\theta}}_n^{(t-1)}, \boldsymbol{\theta}^*\right) \right. \\ &\quad \left. + \frac{1}{2}\left(\widehat{\boldsymbol{\theta}}_n^{(t)} - \boldsymbol{\theta}^*\right)^T \nabla^2_{\boldsymbol{\theta}} h\left(\mathbf{X}_i^{(t-1)}, \widehat{\boldsymbol{\theta}}_n^{(t-1)}, \breve{\boldsymbol{\theta}}_n\left(\mathbf{X}_i^{(t-1)}\right)\right)\left(\widehat{\boldsymbol{\theta}}_n^{(t)} - \boldsymbol{\theta}^*\right)\right) \qquad (37)\end{aligned}$$

where $\breve{\boldsymbol{\theta}}_n\left(\mathbf{X}_i^{(t-1)}\right)$ is an intermediate point between $\widehat{\boldsymbol{\theta}}_n^{(t)}$ and $\boldsymbol{\theta}^*$ (i.e., $\left\|\breve{\boldsymbol{\theta}}_n\left(\mathbf{X}_i^{(t-1)}\right) - \boldsymbol{\theta}^*\right\| \leq \left\|\widehat{\boldsymbol{\theta}}_n^{(t)} - \boldsymbol{\theta}^*\right\|$). The zeroth-order term is $\bar{\mathcal{C}}_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}(\boldsymbol{\theta}^*)$ by definition. To re-express the first-order term, we use the fact, under Assumption 14, that

$$\nabla_{\boldsymbol{\theta}} \bar{\mathcal{C}}_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}\left(\widehat{\boldsymbol{\theta}}_n^{(t)}\right) = 0$$

or equivalently

$$
\begin{aligned}
0 &= \frac{1}{n}\sum_{i=1}^{n} \nabla_{\boldsymbol{\theta}} h\left(\mathbf{X}_i^{(t-1)}, \widehat{\boldsymbol{\theta}}_n^{(t-1)}, \widehat{\boldsymbol{\theta}}_n^{(t)}\right) \\
&= \frac{1}{n}\sum_{i=1}^{n} \nabla_{\boldsymbol{\theta}} h\left(\mathbf{X}_i^{(t-1)}, \widehat{\boldsymbol{\theta}}_n^{(t-1)}, \boldsymbol{\theta}^*\right) + \frac{1}{n}\sum_{i=1}^{n} \nabla^2_{\boldsymbol{\theta}} h\left(\mathbf{X}_i^{(t-1)}, \widehat{\boldsymbol{\theta}}_n^{(t-1)}, \check{\boldsymbol{\theta}}_n\left(\mathbf{X}_i^{(t-1)}\right)\right)\left(\widehat{\boldsymbol{\theta}}_n^{(t)} - \boldsymbol{\theta}^*\right),
\end{aligned}
\tag{38}
$$

where the first-order Taylor expansion is used with an intermediate point $\check{\boldsymbol{\theta}}_n\left(\mathbf{X}_i^{(t-1)}\right)$ between $\widehat{\boldsymbol{\theta}}_n^{(t)}$ and $\boldsymbol{\theta}^*$ (i.e., $\left\|\check{\boldsymbol{\theta}}_n\left(\mathbf{X}_i^{(t-1)}\right) - \boldsymbol{\theta}^*\right\| \leq \left\|\widehat{\boldsymbol{\theta}}_n^{(t)} - \boldsymbol{\theta}^*\right\|$). Rearranging the equation in (38) yields

$$
-\frac{1}{n}\sum_{i=1}^{n} \nabla_{\boldsymbol{\theta}} h\left(\mathbf{X}_i^{(t-1)}, \widehat{\boldsymbol{\theta}}_n^{(t-1)}, \boldsymbol{\theta}^*\right) = \frac{1}{n}\sum_{i=1}^{n} \nabla^2_{\boldsymbol{\theta}} h\left(\mathbf{X}_i^{(t-1)}, \widehat{\boldsymbol{\theta}}_n^{(t-1)}, \check{\boldsymbol{\theta}}_n\left(\mathbf{X}_i^{(t-1)}\right)\right)\left(\widehat{\boldsymbol{\theta}}_n^{(t)} - \boldsymbol{\theta}^*\right).
$$

Plugging this to the equation in (37) results in

$$
\begin{aligned}
\bar{\mathcal{C}}_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}\left(\widehat{\boldsymbol{\theta}}_n^{(t)}\right) &= -\frac{1}{n}\sum_{i=1}^{n}\left(h\left(\mathbf{X}_i^{(t-1)}, \widehat{\boldsymbol{\theta}}_n^{(t-1)}, \boldsymbol{\theta}^*\right) + \left(\widehat{\boldsymbol{\theta}}_n^{(t)} - \boldsymbol{\theta}^*\right)^T \nabla_{\boldsymbol{\theta}} h\left(\mathbf{X}_i^{(t-1)}, \widehat{\boldsymbol{\theta}}_n^{(t-1)}, \boldsymbol{\theta}^*\right)\right. \\
&\quad \left. + \frac{1}{2}\left(\widehat{\boldsymbol{\theta}}_n^{(t)} - \boldsymbol{\theta}^*\right)^T \nabla^2_{\boldsymbol{\theta}} h\left(\mathbf{X}_i^{(t-1)}, \widehat{\boldsymbol{\theta}}_n^{(t-1)}, \check{\boldsymbol{\theta}}_n\left(\mathbf{X}_i^{(t-1)}\right)\right)\left(\widehat{\boldsymbol{\theta}}_n^{(t)} - \boldsymbol{\theta}^*\right)\right) \\
&= \bar{\mathcal{C}}_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}(\boldsymbol{\theta}^*) - \frac{1}{n}\boldsymbol{\delta}_n^T\left(-\frac{1}{n}\sum_{i=1}^{n} \nabla^2_{\boldsymbol{\theta}} h\left(\mathbf{X}_i^{(t-1)}, \widehat{\boldsymbol{\theta}}_n^{(t-1)}, \check{\boldsymbol{\theta}}_n\left(\mathbf{X}_i^{(t-1)}\right)\right)\right)\boldsymbol{\delta}_n \\
&\quad + \frac{1}{2n}\boldsymbol{\delta}_n^T\left(-\frac{1}{n}\sum_{i=1}^{n} \nabla^2_{\boldsymbol{\theta}} h\left(\mathbf{X}_i^{(t-1)}, \widehat{\boldsymbol{\theta}}_n^{(t-1)}, \check{\boldsymbol{\theta}}_n\left(\mathbf{X}_i^{(t-1)}\right)\right)\right)\boldsymbol{\delta}_n.
\end{aligned}
$$

This expression is simplified by defining $\check{\boldsymbol{\Gamma}}_n := -\frac{1}{n}\sum_{i=1}^{n} \nabla^2_{\boldsymbol{\theta}} h\left(\mathbf{X}_i^{(t-1)}, \widehat{\boldsymbol{\theta}}_n^{(t-1)}, \check{\boldsymbol{\theta}}_n\left(\mathbf{X}_i^{(t-1)}\right)\right)$ and $\check{\boldsymbol{\Gamma}}_n := -\frac{1}{n}\sum_{i=1}^{n} \nabla^2_{\boldsymbol{\theta}} h\left(\mathbf{X}_i^{(t-1)}, \widehat{\boldsymbol{\theta}}_n^{(t-1)}, \check{\boldsymbol{\theta}}_n\left(\mathbf{X}_i^{(t-1)}\right)\right)$ as follows:

$$
\bar{\mathcal{C}}_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}\left(\widehat{\boldsymbol{\theta}}_n^{(t)}\right) = \bar{\mathcal{C}}_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}(\boldsymbol{\theta}^*) - \frac{1}{n}\boldsymbol{\delta}_n^T\check{\boldsymbol{\Gamma}}_n\boldsymbol{\delta}_n + \frac{1}{2n}\boldsymbol{\delta}_n^T\check{\boldsymbol{\Gamma}}_n\boldsymbol{\delta}_n.
$$

We note that for the first term, $\mathbb{E}\left[\bar{\mathcal{C}}_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}(\boldsymbol{\theta}^*)\right] = \mathcal{C}(\boldsymbol{\theta}^*)$. For the second term, the mean value theorem applied to the element at the k^{th} row and l^{th} column of the i^{th}

summand in $\check{\mathbf{\Gamma}}_n$ yields

$$\begin{aligned}\left[\nabla^2_{\boldsymbol{\theta}} h\left(\mathbf{X}_i^{(t-1)}, \widehat{\boldsymbol{\theta}}_n^{(t-1)}, \check{\boldsymbol{\theta}}_n\left(\mathbf{X}_i^{(t-1)}\right)\right)\right]_{kl} &= \left[\nabla^2_{\boldsymbol{\theta}} h\left(\mathbf{X}_i^{(t-1)}, \boldsymbol{\theta}^*, \boldsymbol{\theta}^*\right)\right]_{kl} \\ &+ \left(\widehat{\boldsymbol{\theta}}_n^{(t-1)} - \boldsymbol{\theta}^*\right)^T \nabla_{\boldsymbol{\eta}}\left[\nabla^2_{\boldsymbol{\theta}} h\left(\mathbf{X}_i^{(t-1)}, \widehat{\boldsymbol{\theta}}'_n, \check{\boldsymbol{\theta}}'_n\right)\right]_{kl} \\ &+ \left(\check{\boldsymbol{\theta}}_n\left(\mathbf{X}_i^{(t-1)}\right) - \boldsymbol{\theta}^*\right)^T \nabla_{\boldsymbol{\theta}}\left[\nabla^2_{\boldsymbol{\theta}} h\left(\mathbf{X}_i^{(t-1)}, \widehat{\boldsymbol{\theta}}'_n, \check{\boldsymbol{\theta}}'_n\right)\right]_{kl},\end{aligned}$$

where $\widehat{\boldsymbol{\theta}}'_n$ is an intermediate value between $\widehat{\boldsymbol{\theta}}_n^{(t-1)}$ and $\boldsymbol{\theta}^*$. Similarly, $\check{\boldsymbol{\theta}}'_n$ is an intermediate value between $\check{\boldsymbol{\theta}}_n\left(\mathbf{X}_i^{(t-1)}\right)$ and $\boldsymbol{\theta}^*$. Under Assumption 15, the last two terms converge to zero in probability as $n \to \infty$ so that the left-hand side converges to the first term of the right-hand side in probability. The average over i yields that $\check{\mathbf{\Gamma}}_n$ converges in probability to $\mathbf{\Gamma}$ as $n \to \infty$ by the weak law of large numbers. Similarly, $\check{\mathbf{\Gamma}}_n$ converges in probability to $\mathbf{\Gamma}$.

Similar to the derivations leading to (36), both $\boldsymbol{\delta}_n^T \check{\mathbf{\Gamma}}_n \boldsymbol{\delta}_n$ and $\boldsymbol{\delta}_n^T \check{\mathbf{\Gamma}}_n \boldsymbol{\delta}_n$ converge in distribution to $\rho\chi^2_d$ by Slutsky's theorem so that their expectations converge to ρd as $n \to \infty$ under Assumption 16 that imposes uniform integrability on both $\boldsymbol{\delta}_n^T \check{\mathbf{\Gamma}}_n \boldsymbol{\delta}_n$ and $\boldsymbol{\delta}_n^T \check{\mathbf{\Gamma}}_n \boldsymbol{\delta}_n$. Thus,

$$\mathbb{E}\left[\bar{\mathcal{C}}_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}\left(\widehat{\boldsymbol{\theta}}_n^{(t)}\right)\right] = \mathcal{C}(\boldsymbol{\theta}^*) - \frac{\rho d}{n} + \frac{\rho d}{2n} + o\left(\frac{1}{n}\right). \tag{39}$$

Therefore, combining (36) and (39), the bias of interest is

$$\begin{aligned}\mathbb{E}\left[\bar{\mathcal{C}}_{\widehat{\boldsymbol{\theta}}_n^{(t-1)}}\left(\widehat{\boldsymbol{\theta}}_n^{(t)}\right) - \mathcal{C}\left(\widehat{\boldsymbol{\theta}}_n^{(t)}\right)\right] &= \mathcal{C}(\boldsymbol{\theta}^*) - \frac{\rho d}{n} + \frac{\rho d}{2n} - \left(\mathcal{C}(\boldsymbol{\theta}^*) + \frac{\rho d}{2n}\right) + o\left(\frac{1}{n}\right) \\ &= -\frac{\rho d}{n} + o\left(\frac{1}{n}\right).\end{aligned}$$

□

Appendix B: Implementation Details

For the implementation of the EM algorithm, we use multiple random initial values of $\boldsymbol{\theta}$ and choose the best minimizer of $\bar{\mathcal{C}}^{(t-1)}(\boldsymbol{\theta})$ in (17) (Figueiredo and Jain, 2002) to reduce the impact of initial guess of $\boldsymbol{\theta}$ on the algorithm's performance and avoid getting

stuck with a local minimizer. In the numerical study in Section 5, we choose the initial $\boldsymbol{\theta}$'s by drawing $\mu_1, \ldots, \mu_{30}$ from a standard multivariate Gaussian and setting all $\Sigma_1, \ldots, \Sigma_{30}$ as $3I_{p\times p}$ at the 0^{th} iteration. At the t^{th} iteration for $t \geq 1$, we randomly select $\mu_1, \ldots, \mu_k$ from $\{\mathbf{X}_i^{(s-1)} : h\Big(\mathbf{X}_i^{(s-1)}, \widehat{\boldsymbol{\theta}}_n^{(s-1)}, \boldsymbol{\theta}\Big) > 0, i = 1, \ldots, n; s = 1, \ldots, t\}$ without replacement. However, if the set's cardinality is smaller than k, we randomly select any elements in $\{\mathbf{X}_i^{(s-1)} : h\Big(\mathbf{X}_i^{(s-1)}, \widehat{\boldsymbol{\theta}}_n^{(s-1)}, \boldsymbol{\theta}\Big) = 0, i = 1, \ldots, n; s = 1, \ldots, t\}$ for the remaining parameters. We set $\Sigma_1, \ldots, \Sigma_k$ as $(3/p)\,\text{trace}\big(cov(\bar{\mathbf{X}})\big)\, I_{p\times p}$, where $\bar{\mathbf{X}}$ is the data matrix created by augmenting $\{\mathbf{X}_i^{(s-1)} : i = 1, \ldots, n; s = 1, \ldots, t\}$, and *cov* is the sample covariance.

When the number of components k is set large enough to cause an overfitting issue within the EM algorithm, we catch it by monitoring the condition numbers of the Gaussian components' covariances (Figueiredo and Jain, 2002). In the numerical study in Section 5, we abort the EM algorithm when the condition number of any covariance exceeds 10^5. If we need to abort most of the EM algorithms that start with different initial parameter guesses (we use the threshold of 5 aborted out of 10), it indicates that k is already too large for the given sample size.

To check the convergence of the EM algorithm, we check the reduction of $\bar{\mathcal{C}}^{(t-1)}(\boldsymbol{\theta})$ in (17). In the numerical study in Section 5, we stop iterating updating equations in the EM algorithm if the reduction of $\bar{\mathcal{C}}^{(t-1)}(\boldsymbol{\theta})$ is less than 1% or a specified maximum number of iterations is reached.

We compute the CIC for $k = k_{\min}, k_{\min}+1, \ldots, k_{\max}$ for the grid search of the minimizer $k^{*(t)}$ at the t^{th} iteration. We start from $k = k_{\min}$ (one, typically) and increase k. In practice, it is generally unnecessary to go up to $k_{\max}$, which is upper bounded by a known function of the sample size, because the overfitting is detected within the EM algorithm. To reduce the grid search time, we compute the moving average of the CIC and stop increasing k when the moving average starts to increase. In our numerical study in Section 5, we use the window of four k values for the moving average.

Acknowledgment

The author thanks Dr. Yen-Chi Chen (University of Washington) for reviewing an earlier version of this paper and providing detailed feedback, and Dr. Archis Ghate (University of Washington) for suggesting the reduction of the CIC to the AIC. This work was conceived during the author's Ph.D. study supervised by Dr. Eunshin Byon (University of Michigan) who first informed the author of the model order selection problem in mixture model-based importance sampling.

References

H. Akaike. Information theory and the maximum likelihood principle. In B. N. Petrov and F. Czáki, editors, *2nd International Symposium on Information Theory*, pages 267–281. Budapest: Akademiai Kiadó, 1973.

H. Akaike. A new look at the statistical model identification. *IEEE Transactions on Automatic Control*, 19(6):716–723, 1974.

F. Bassetti and P. Diaconis. Examples comparing importance sampling and the metropolis algorithm. *Illinois Journal of Mathematics*, 50(1-4):67–91, 2006.

M. A. Beaumont. Approximate Bayesian computation in evolution and ecology. *Annual Review of Ecology, Evolution, and Systematics*, 41:379–406, 2010.

D. M. Blei, A. Kucukelbir, and J. D. McAuliffe. Variational inference: A review for statisticians. *arXiv preprint arXiv:1601.00670*, 2016.

Z. I. Botev, D. P. Kroese, R. Y. Rubinstein, and P. L'Ecuyer. The cross-entropy method for optimization. *Machine Learning: Theory and Applications, V. Govindaraju and C. R. Rao, Eds, Chennai: Elsevier*, 31:35–59, 2013.

K. P. Burnham and D. R. Anderson. *Model selection and multimodel inference: a practical information-theoretic approach*. New York: Springer Science & Business Media, 2003.

Y. Choe, E. Byon, and N. Chen. Importance sampling for reliability evaluation with stochastic simulation models. *Technometrics*, 57(3):351–361, 2015.

A. P. Dempster, N. M. Laird, and D. B. Rubin. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 1–38, 1977.

M. C. Donohue, R. Overholser, R. Xu, and F. Vaida. Conditional Akaike information under generalized linear and proportional hazards mixed models. *Biometrika*, 98 (3):685–700, 2011.

M. A. Figueiredo and A. K. Jain. Unsupervised learning of finite mixture models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24(3):381–396, 2002.

A. E. Gelfand and A. F. Smith. Sampling-based approaches to calculating marginal densities. *Journal of the American Statistical Association*, 85(410):398–409, 1990.

G. H. Givens and A. E. Raftery. Local adaptive importance sampling for multivariate densities with strong nonlinear relationships. *Journal of the American Statistical Association*, 91(433):132–141, 1996.

W. K. Hastings. Monte Carlo sampling methods using Markov chains and their applications. *Biometrika*, 57(1):97–109, 1970.

T. C. Hesterberg. Weighted average importance sampling and defensive mixture distributions. *Technometrics*, 37(2):185–194, 1995.

P. J. Huber. Robust estimation of a location parameter. *The Annals of Mathematical Statistics*, 35(1):73–101, 1964.

M. I. Jordan, Z. Ghahramani, T. S. Jaakkola, and L. K. Saul. An introduction to variational methods for graphical models. *Machine Learning*, 37(2):183–233, 1999.

H. Kahn and A. W. Marshall. Methods of reducing sample size in Monte Carlo computations. *Journal of the Operations Research Society of America*, 1(5):263–278, 1953.

K. H. Knuth, M. Habeck, N. K. Malakar, A. M. Mubeen, and B. Placek. Bayesian evidence and model selection. *Digital Signal Processing*, 47:50–67, 2015.

S. Kullback and R. A. Leibler. On information and sufficiency. *The Annals of Mathematical Statistics*, 22(1):79–86, 1951.

N. Kurtz and J. Song. Cross-entropy-based adaptive importance sampling using gaussian mixture. *Structural Safety*, 42:35–44, 2013.

J. S. Liu. Metropolized independent sampling with comparisons to rejection sampling and importance sampling. *Statistics and Computing*, 6(2):113–119, 1996.

N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller. Equation of state calculations by fast computing machines. *The Journal of Chemical Physics*, 21(6):1087–1092, 1953.

J. C. Neddermeyer. Computationally efficient nonparametric importance sampling. *Journal of the American Statistical Association*, 104(486):788–802, 2009.

A. Owen and Y. Zhou. Safe and effective importance sampling. *Journal of the American Statistical Association*, 95(449):135–143, 2000.

R. Rubinstein. The cross-entropy method for combinatorial and continuous optimization. *Methodology and Computing in Applied Probability*, 1(2):127–190, 1999.

R. Y. Rubinstein and A. Shapiro. *Discrete event systems: Sensitivity analysis and stochastic optimization by the score function method*. Chichester: John Wiley & Sons Ltd, 1993.

G. Schwarz. Estimating the dimension of a model. *The Annals of Statistics*, 6(2): 461–464, 1978.

M. Sunnåker, A. G. Busetto, E. Numminen, J. Corander, M. Foll, and C. Dessimoz. Approximate Bayesian computation. *PLOS Computational Biology*, 9(1):1–10, 01 2013.

A. W. Van der Vaart. *Asymptotic statistics*. New York: Cambridge University Press, 1998.

M. J. Wainwright and M. I. Jordan. Graphical models, exponential families, and variational inference. *Foundations and Trends® in Machine Learning*, 1(1–2):1–305, 2008.

H. Wang and X. Zhou. A cross-entropy scheme for mixtures. *ACM Transactions on Modeling and Computer Simulation*, 25(1):6:1–6:20, 2015. ISSN 1049-3301.

P. Zhang. Nonparametric importance sampling. *Journal of the American Statistical Association*, 91(435):1245–1253, 1996.